CRAY XT3
Measured Balance

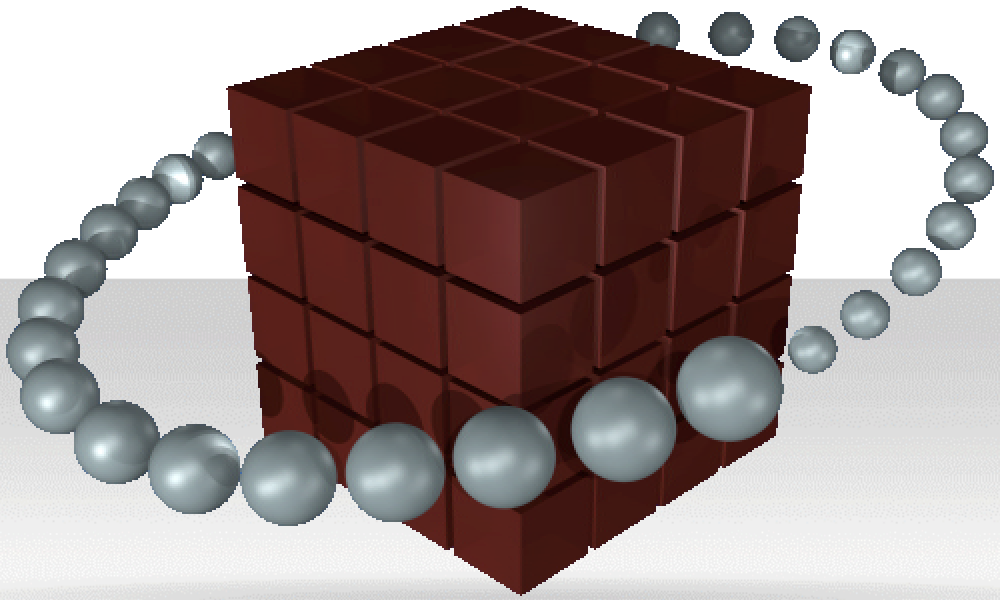

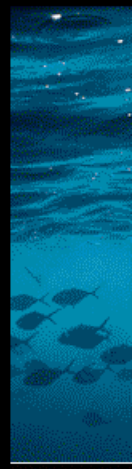

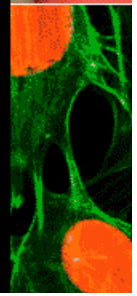

Cray XT3 Processing Element: Measured Performance

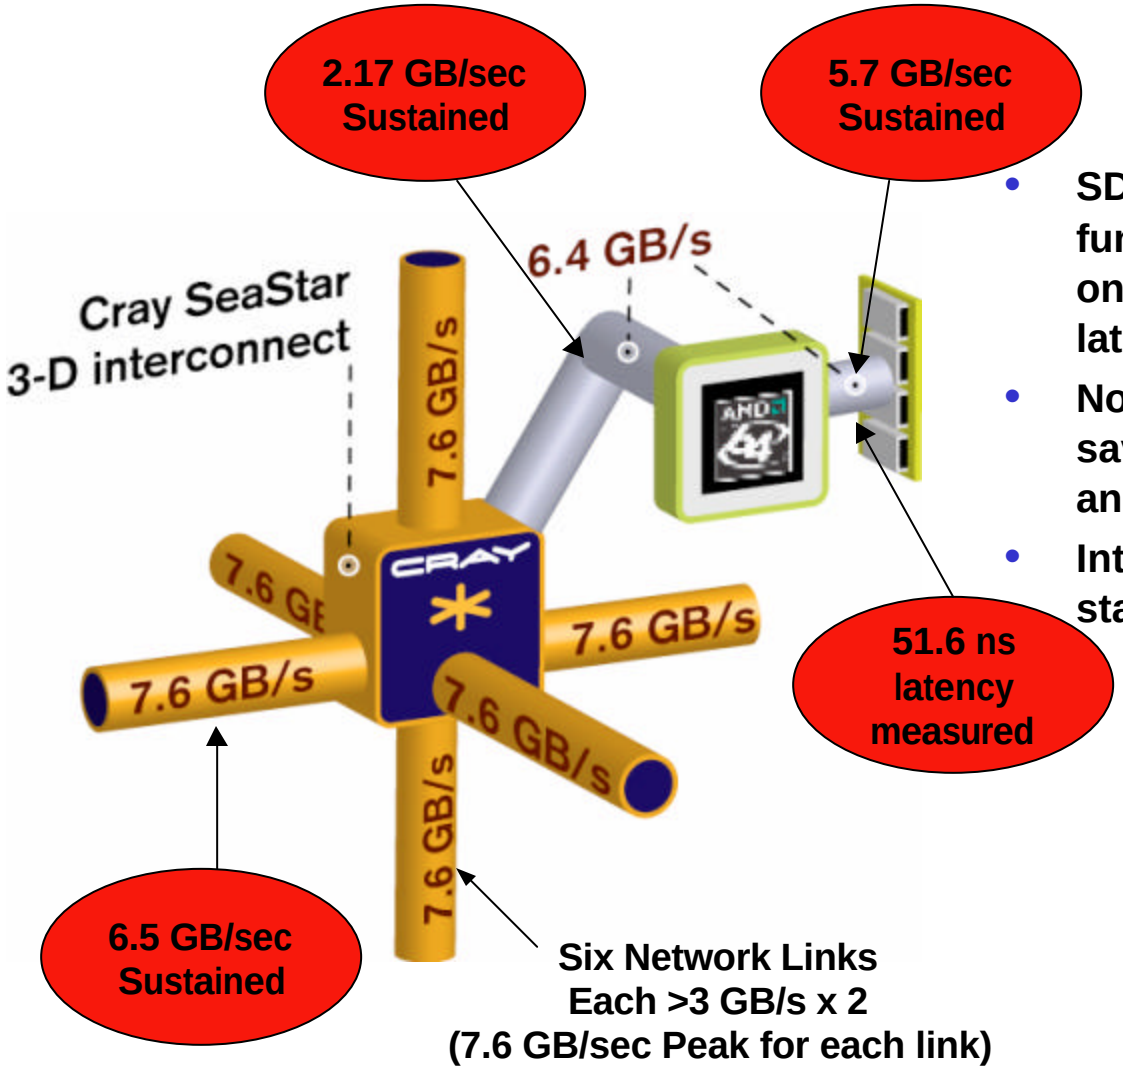

- **SDRAM memory controller and function of Northbridge is pulled onto the Opteron die. Memory latency reduced to <60 ns**
- **No Northbridge chip results in savings in heat, power, complexity and an increase in performance**
- **Interface off the chip is an open standard (HyperTransport)**

Memory Latency

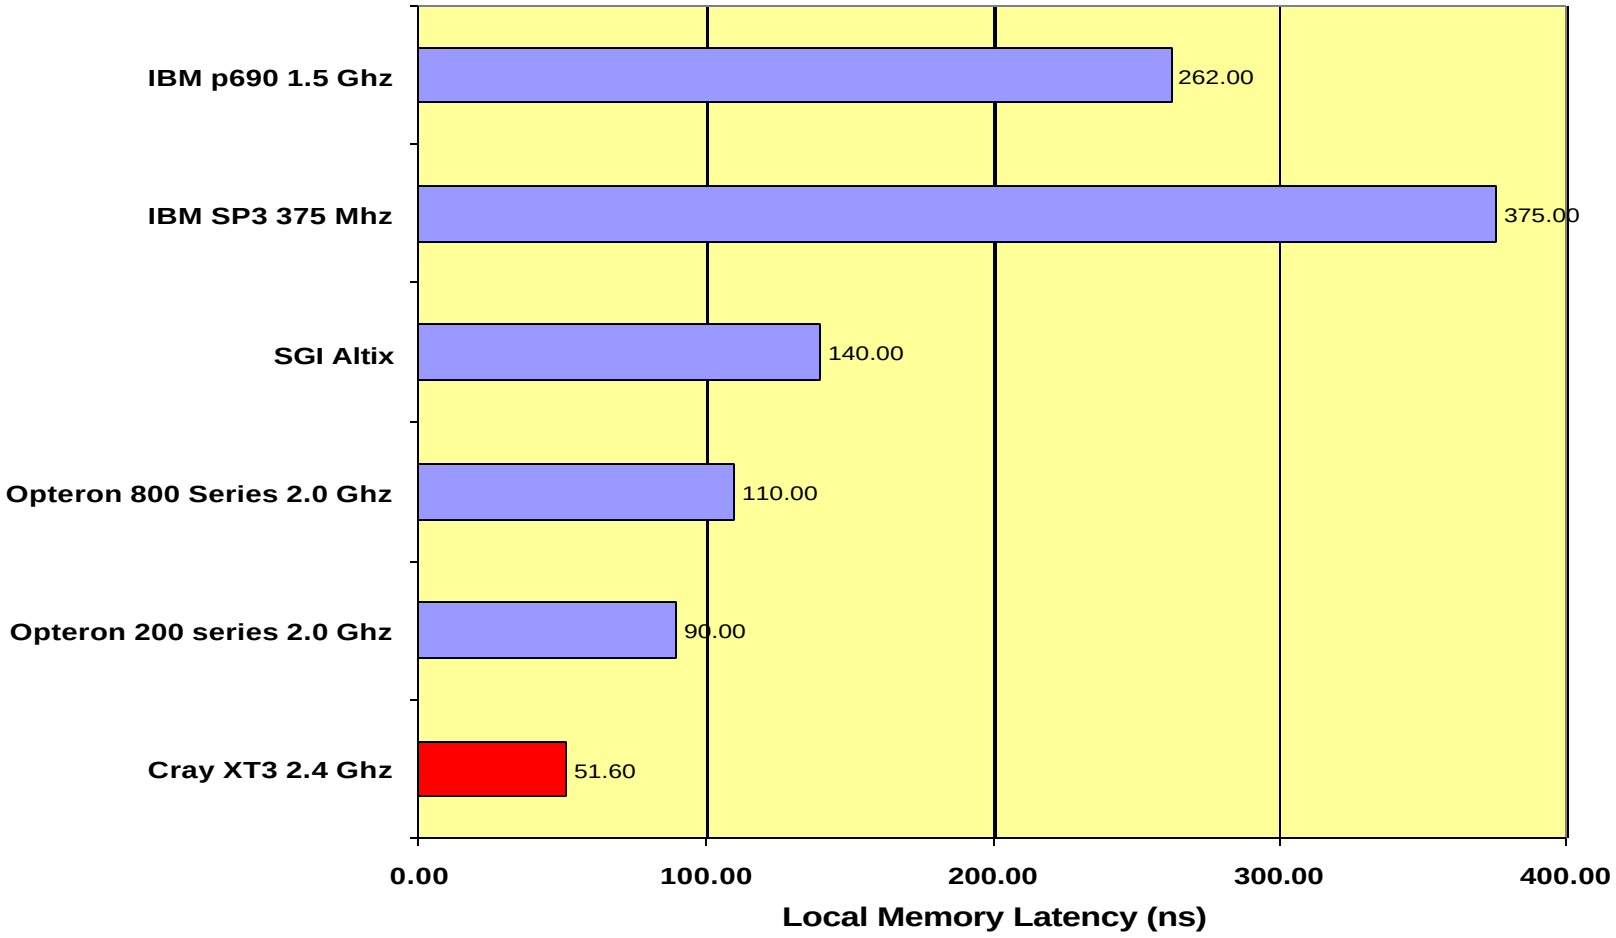

Single Processor architecture yields lowest memory latency

HPCC Random Access Benchmark

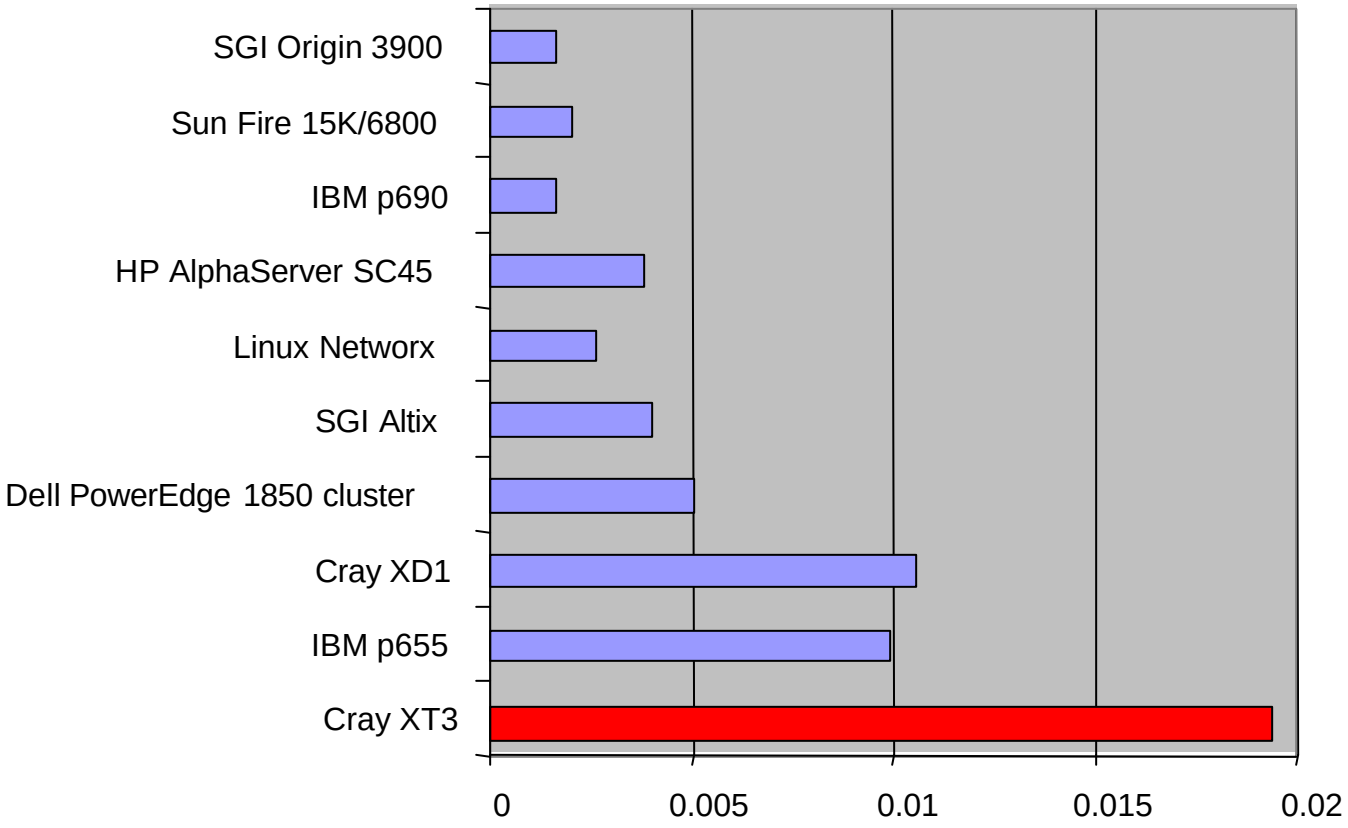

Measured Memory Balance

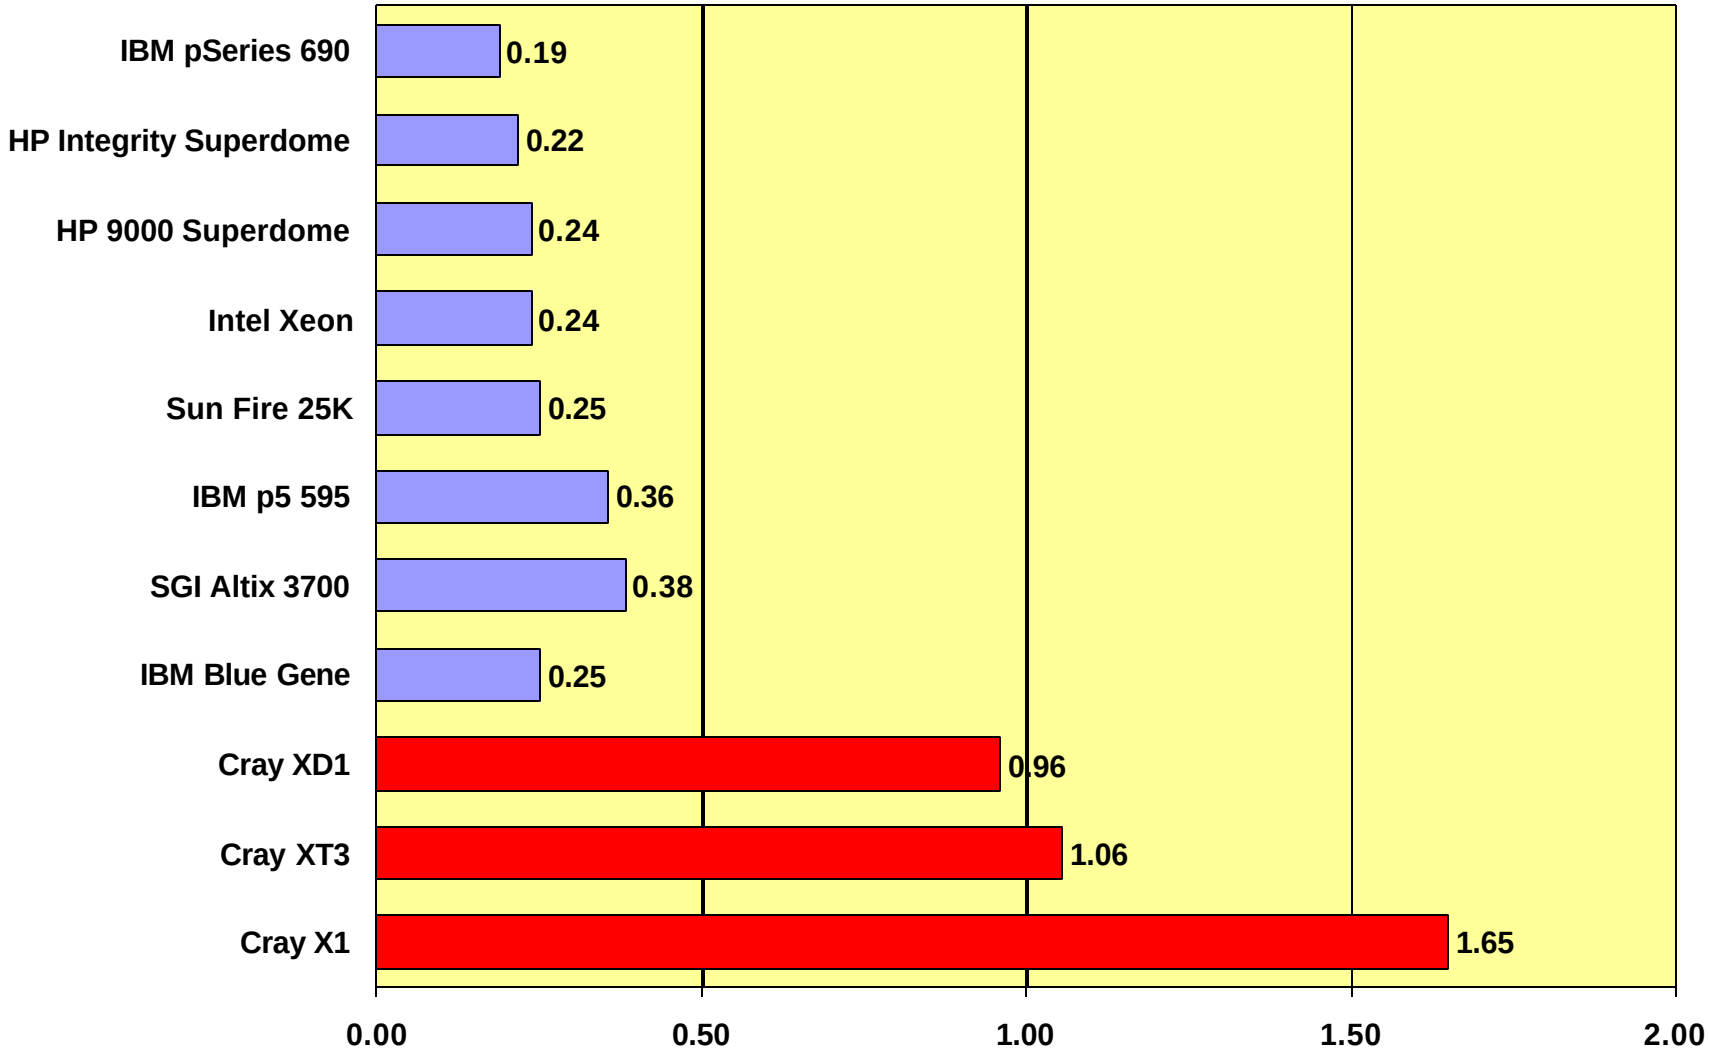

B/F calculated from memory bandwidth measured via STREAM Triad benchmark

Measured Network Balance

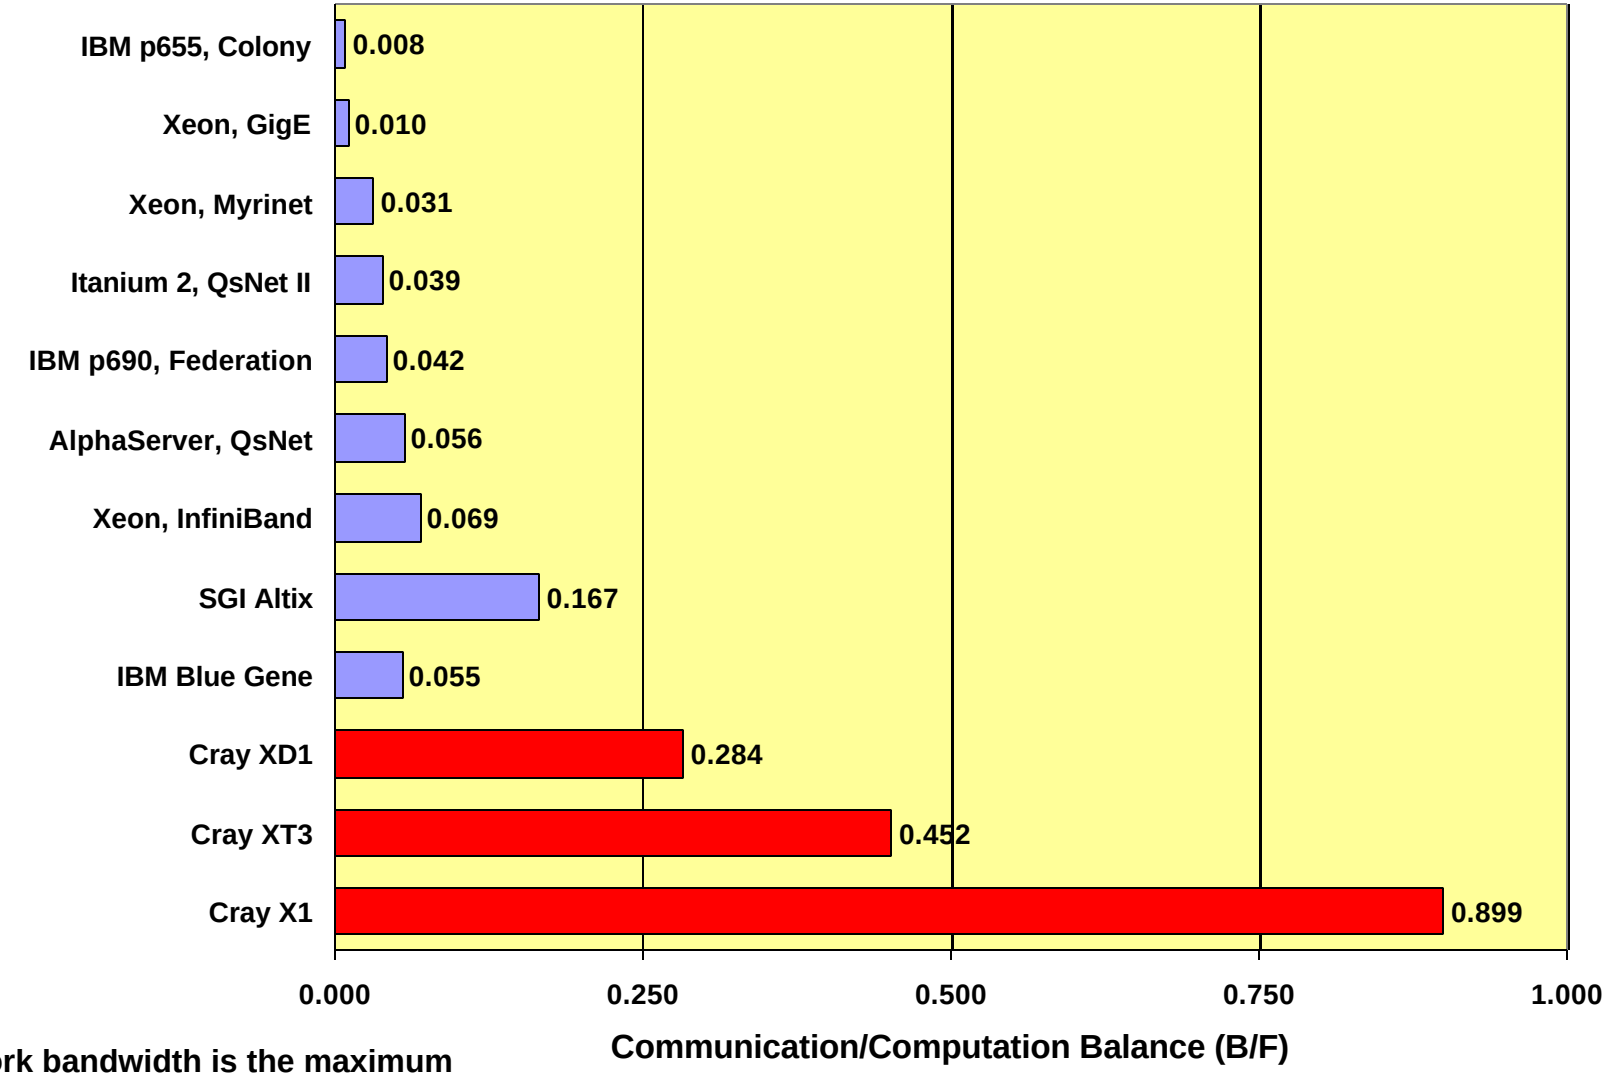

Network bandwidth is the maximum bidirectional data exchange rate between two nodes using MPI

Cray SeaStar Internals

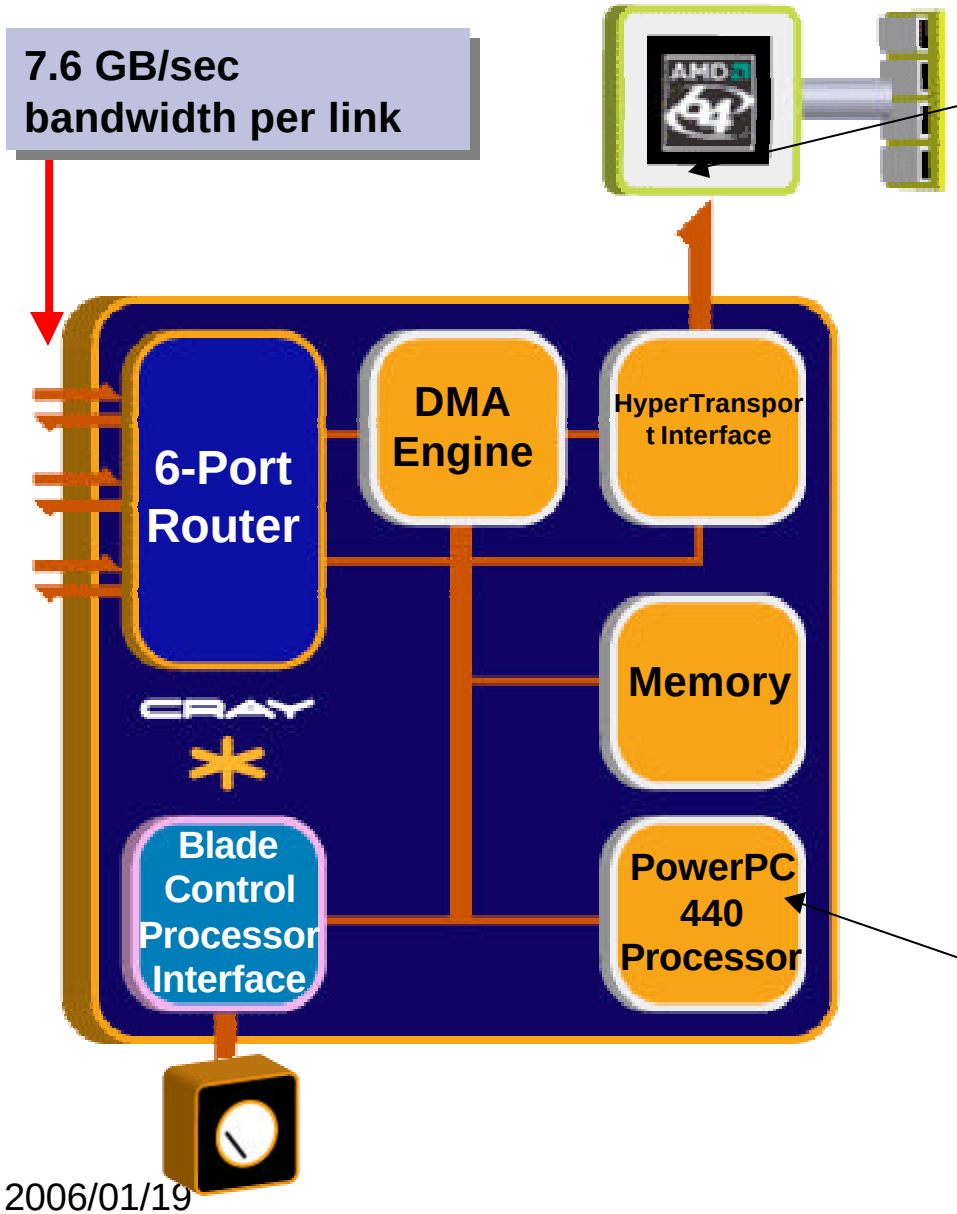

Generic over Integrated. Linux & Catamount functionality

- Each Processor is directly connected to a dedicated SeaStar
- Each SeaStar contains a 6-Port router *and* communications engine
- Provides serial connection to the Cray RAS and Management Sy

Accelerated Portals for Catamount June 05 Performance

The Supercomputer Company

Cray XT3
Early Results

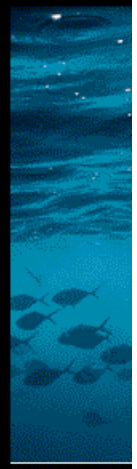


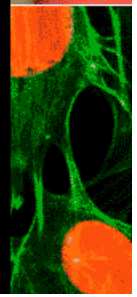

Sandia 7X Applications Progress:

Code	Version	1/24/05	3/9/05	3/20/05
Alegra	4.5	8 PEs	256 PEs	256 PEs
CTH	6.0 Mar 02	44 PEs	1024 PEs	2048 PEs
ITS	6/23/04	16 PEs	670 PEs	3827 PEs
SAGE	3/10/03	16 PEs	1024 PEs	1900 PEs
Partisn	9/16/02	44 PEs	564 PEs	1900 PEs
UMT2000	1.2.2 (1/28/02)	4 PEs	1024 PEs	3000 PEs
sPPM	2.0 (1/1/04)	16 PEs	1331 PEs	3600 PEs
Salinas	1.2	ported	64 PEs	1000 PEs
Presto		Compile Issues	256 PEs	1536 PEs
Calore		Compile Issues	256 PEs	1024 PEs

Stream Benchmark

Function	T3E1200E (MB/sec)	CRAY XT3 (MB/sec)	Ratio
Copy:	**520**	**5755**	**11.1**
Scale:	**517**	**4464**	**8.6**
Add:	**611**	**4142**	**6.8**
Triad:	**622**	**5549**	**8.9**

Measured on a 2.4 Ghz Opteron with PC3200 DDR DIMMS. Tuned assembler code

Stream Benchmark

Function	PGI 6.0 (MB/sec)	Assembler (MB/sec)
Copy:	**4930**	**5755**
Scale:	**4930**	**4464**
Add:	**5030**	**4142**
Triad:	**5070**	**5549**

The pre-fetch option approximate what the assembler code does…

NAS Kernels

- **All results in Mflops/second (64-bit)**
- **No changes to source code**

Kernel	Cray T3E900	CRAY XT3 2.4Ghz	Speedup
MXM	174	2223	12.7
CFFT2D	23	775	33.7
CHOLSKY	26	578	22.2
BTRIX	48	1291	26.9
GMTRY	73	472	6.5
EMIT	246	956	3.9
VPENTA	26	146	5.6
Average			15.9

Interconnect Performance

- Full N x N network run on two cabinets
- Network Topology was 2 x 4 x 24
- Bi-Section Bandwidth across a 2 x 4 “plane” measured at 52.5 GB/sec
- This nets out to 6.5 GB/sec payload bandwidth per link

FFTW Performance

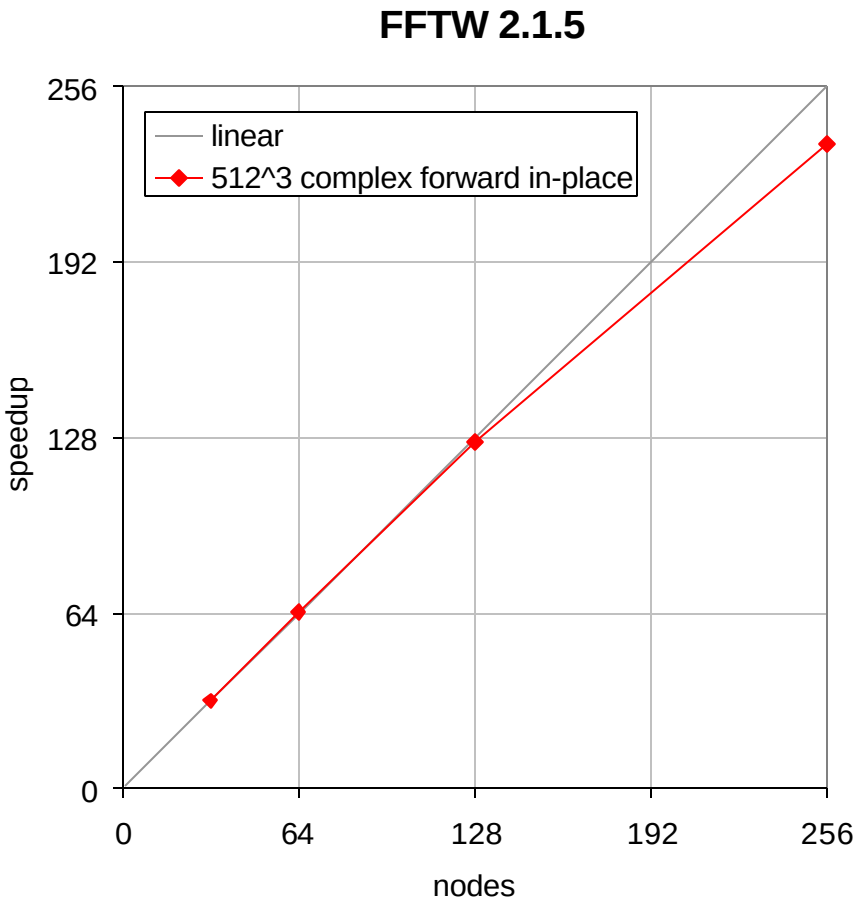

- **Favorable scaling on FFTs and other transpose-intensive operations is essential to numerous applications**

nodes	efficiency
32	1
64	1.01
128	0.990
256	0.918

Standard Benchmarks

- *Performance numbers are extremely preliminary and will improve as the system matures*
- HPCC (1848 nodes, 2.0 Ghz)
  - HPL : 4,079 GFlop/s (55% of theoretical peak on very small problem run)
  - PTRANS : 217 GB/s (not tuned, no torus)
- HPCC (128 nodes 2.4 Ghz)
  - EP DGEMM : 4.26 GFlop/s per processor
  - EP GUPS : .019 billion updates / sec per processor
- Pallas MPI Benchmarks
  - ping-pong bandwidth: 1094 MB/s
  - Send-receive benchmark: 2170 MB/sec

CG: Class C

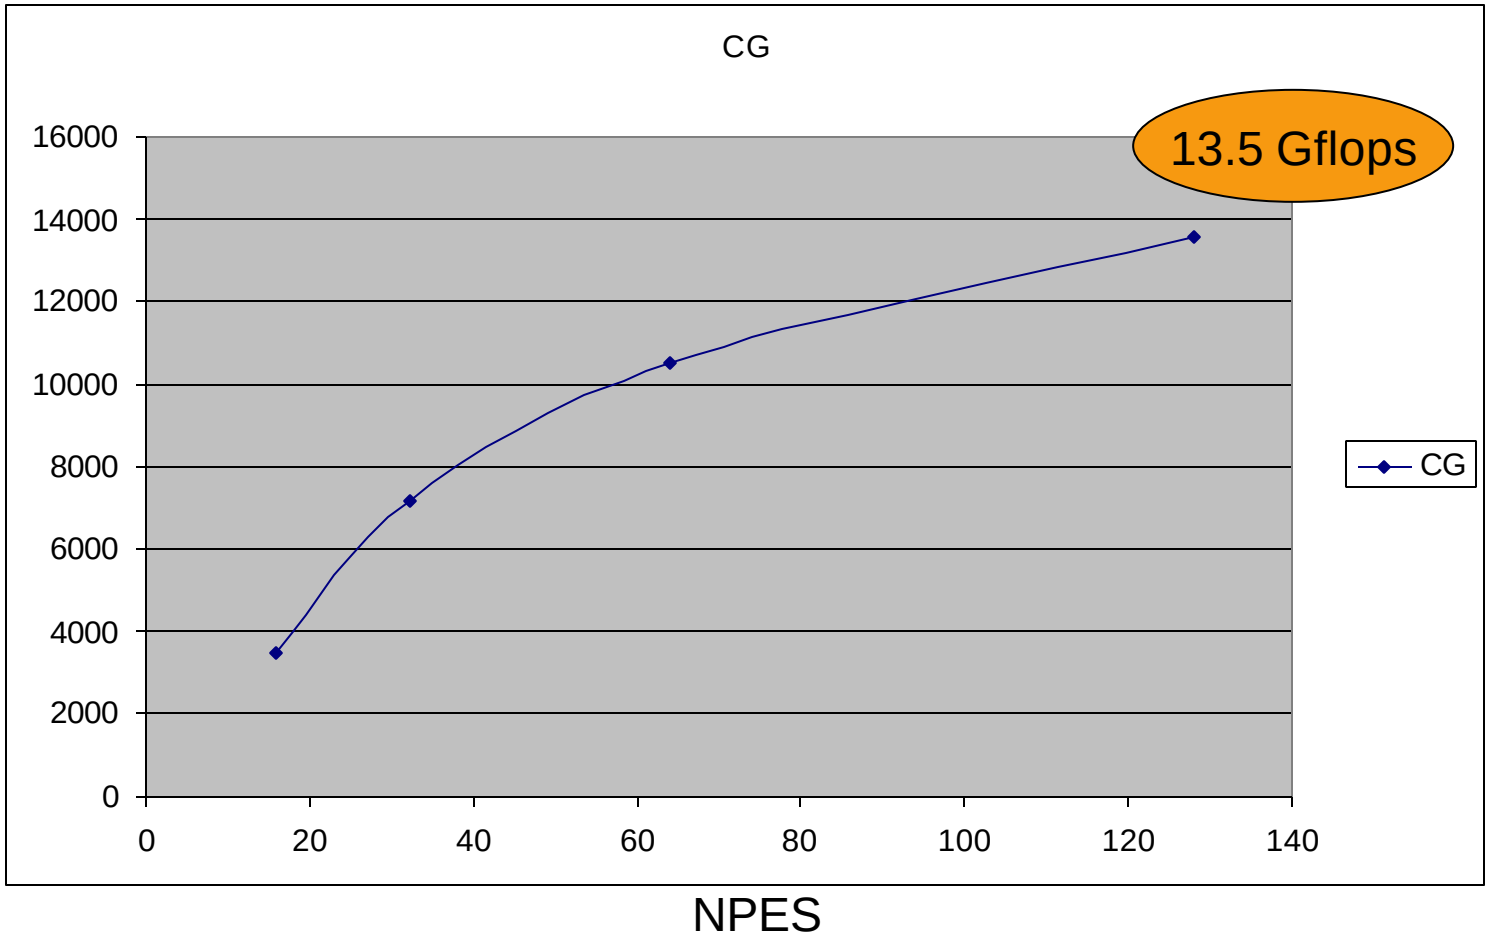
CG
13.5 Gflops
16000
14000
12000
10000
8000
6000
4000
2000
0
0
20
40
60
80
100
120
140
CG
Mflops
NPES

EP: Class C

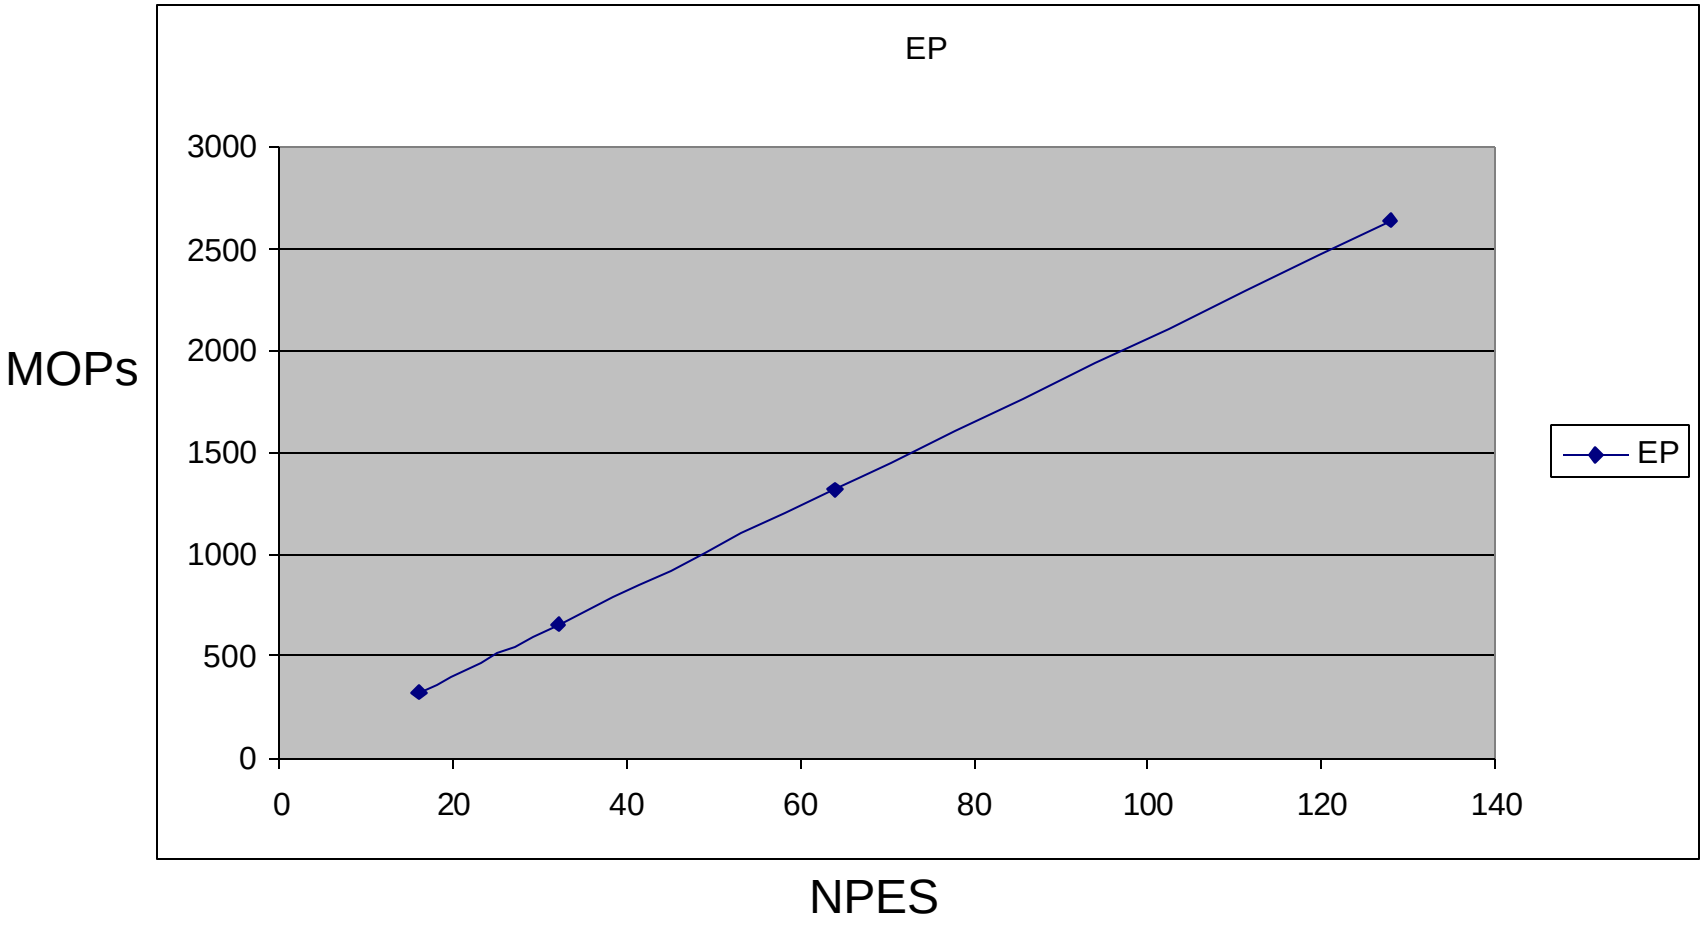
EP
MOPs
3000
2500
2000
1500
1000
500
0
0
20
40
60
80
100
120
140
NPES
EP

MG: Class C

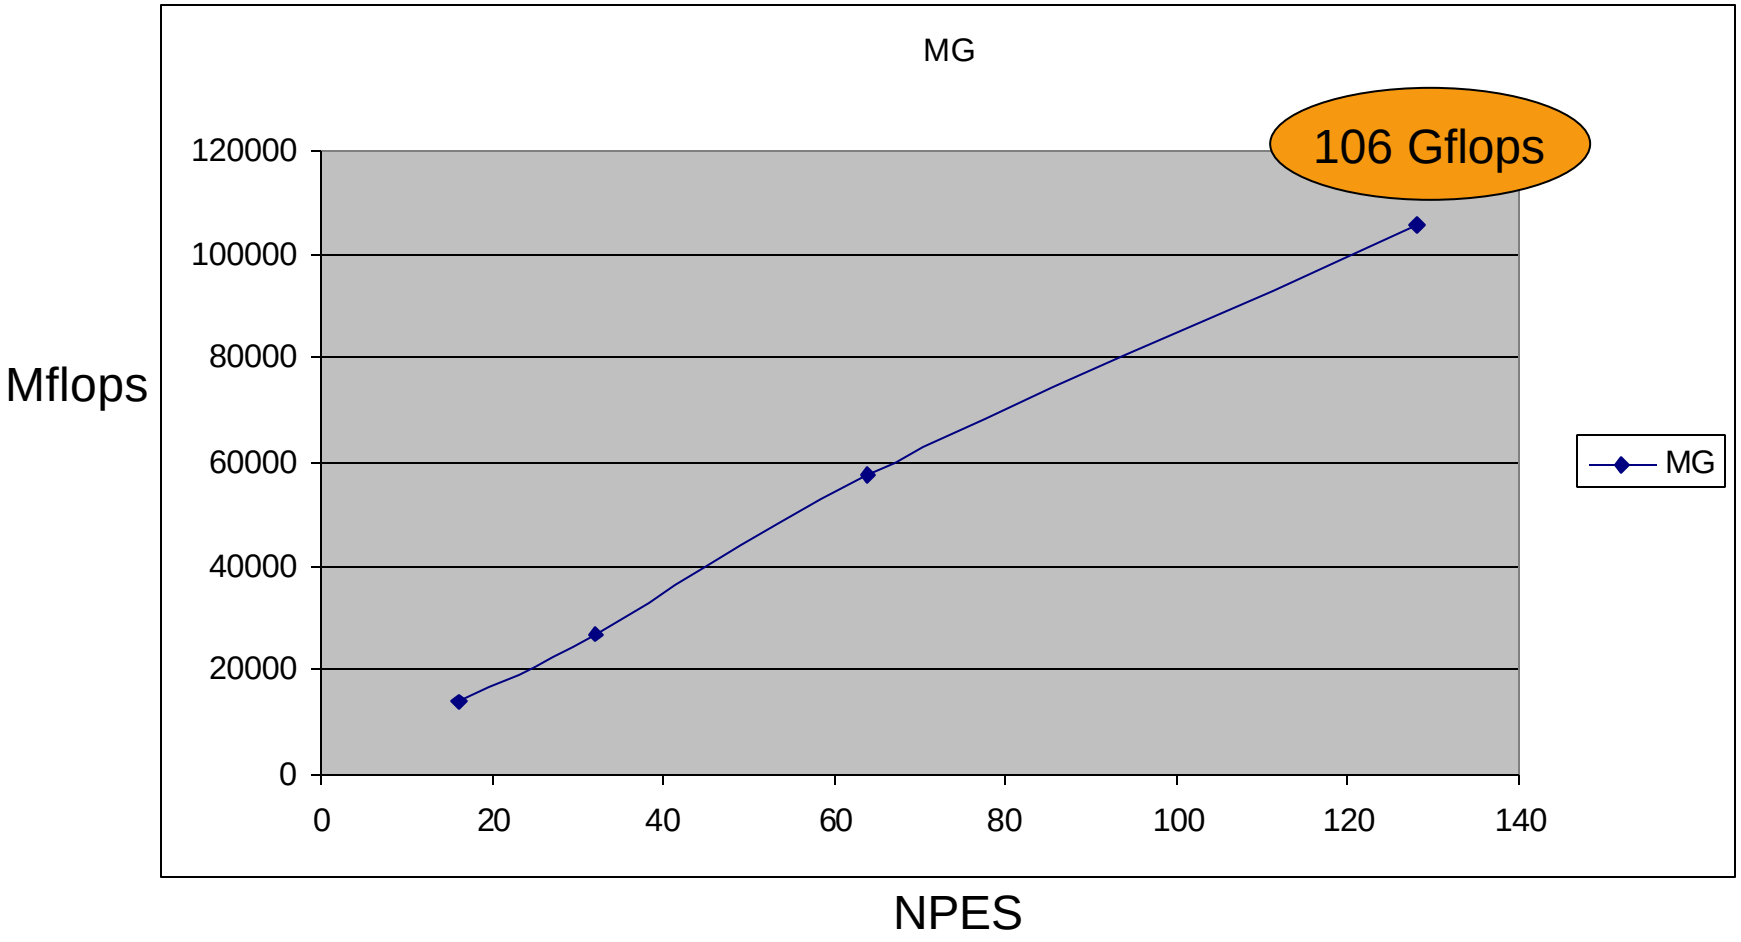
MG
106 Gflops
120000
100000
80000
60000
40000
20000
0
Mflops
0
20
40
60
80
100
120
140
NPES
MG

FT: Class C

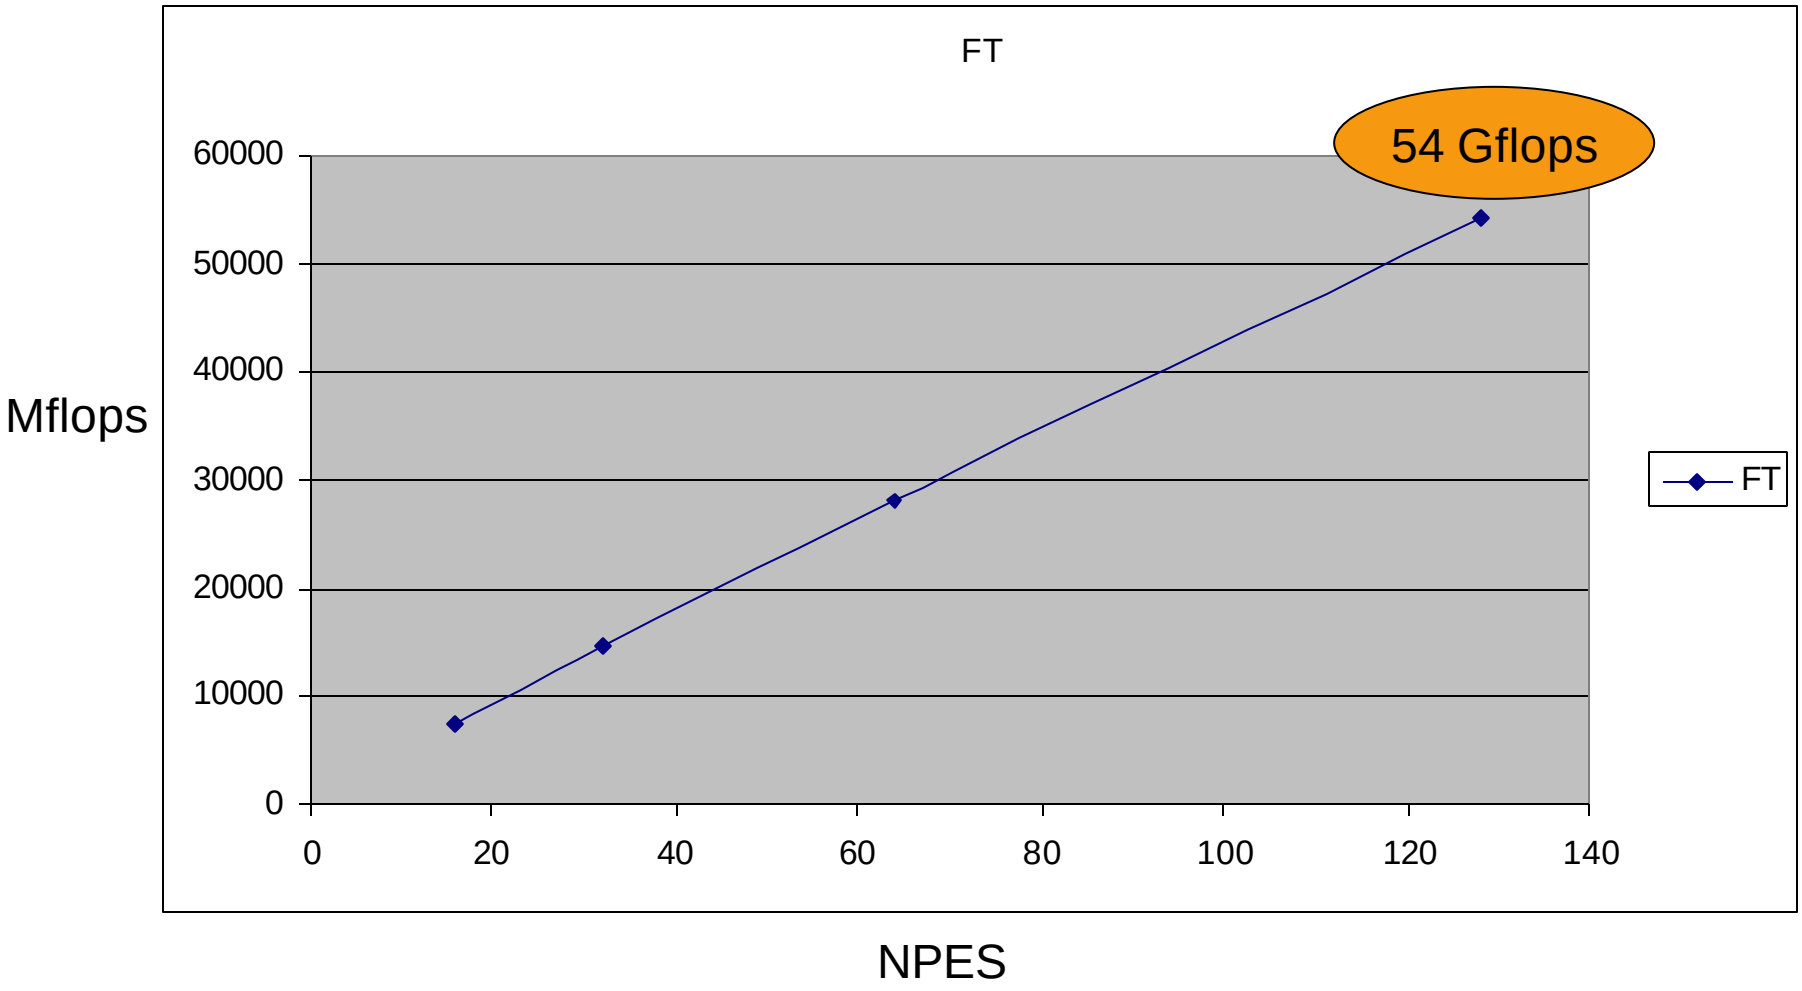
FT
54 Gflops
60000
50000
40000
30000
20000
10000
0
Mflops
0
20
40
60
80
100
120
140
NPES
FT

IS: Class C

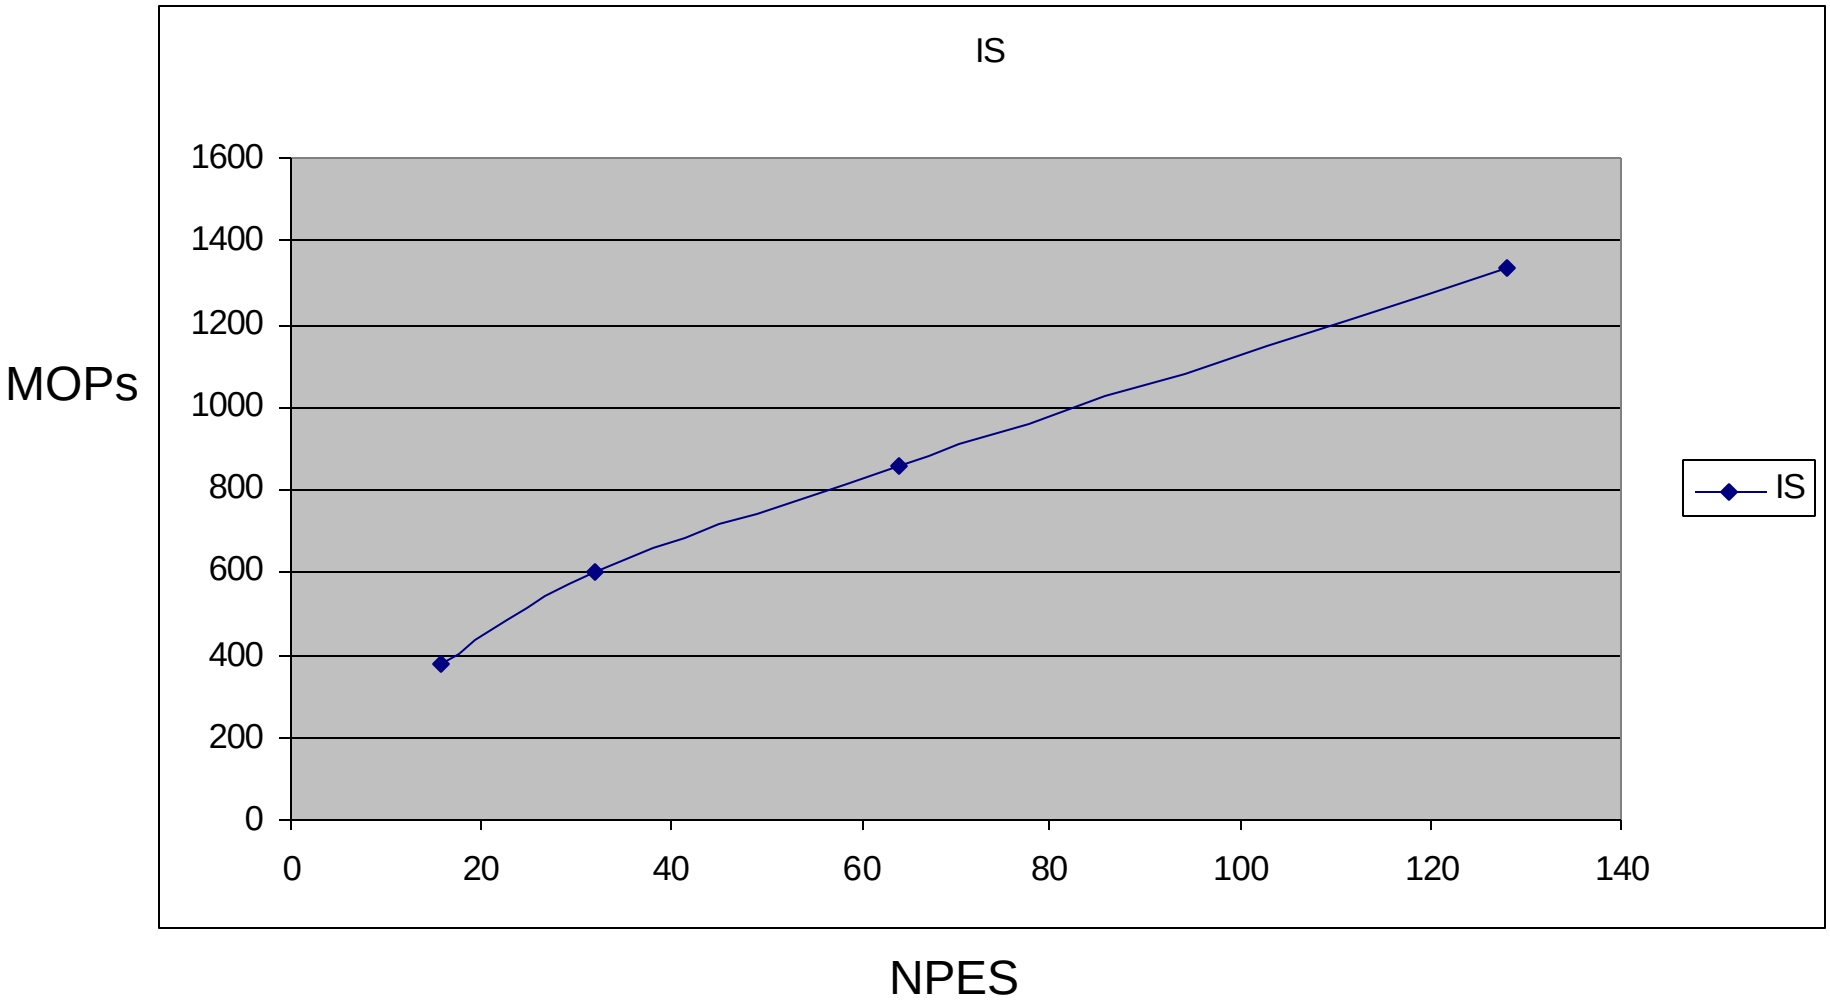
IS
MOPs
1600
1400
1200
1000
800
600
400
200
0
0
20
40
60
80
100
120
140
IS
NPES

LU: Class C

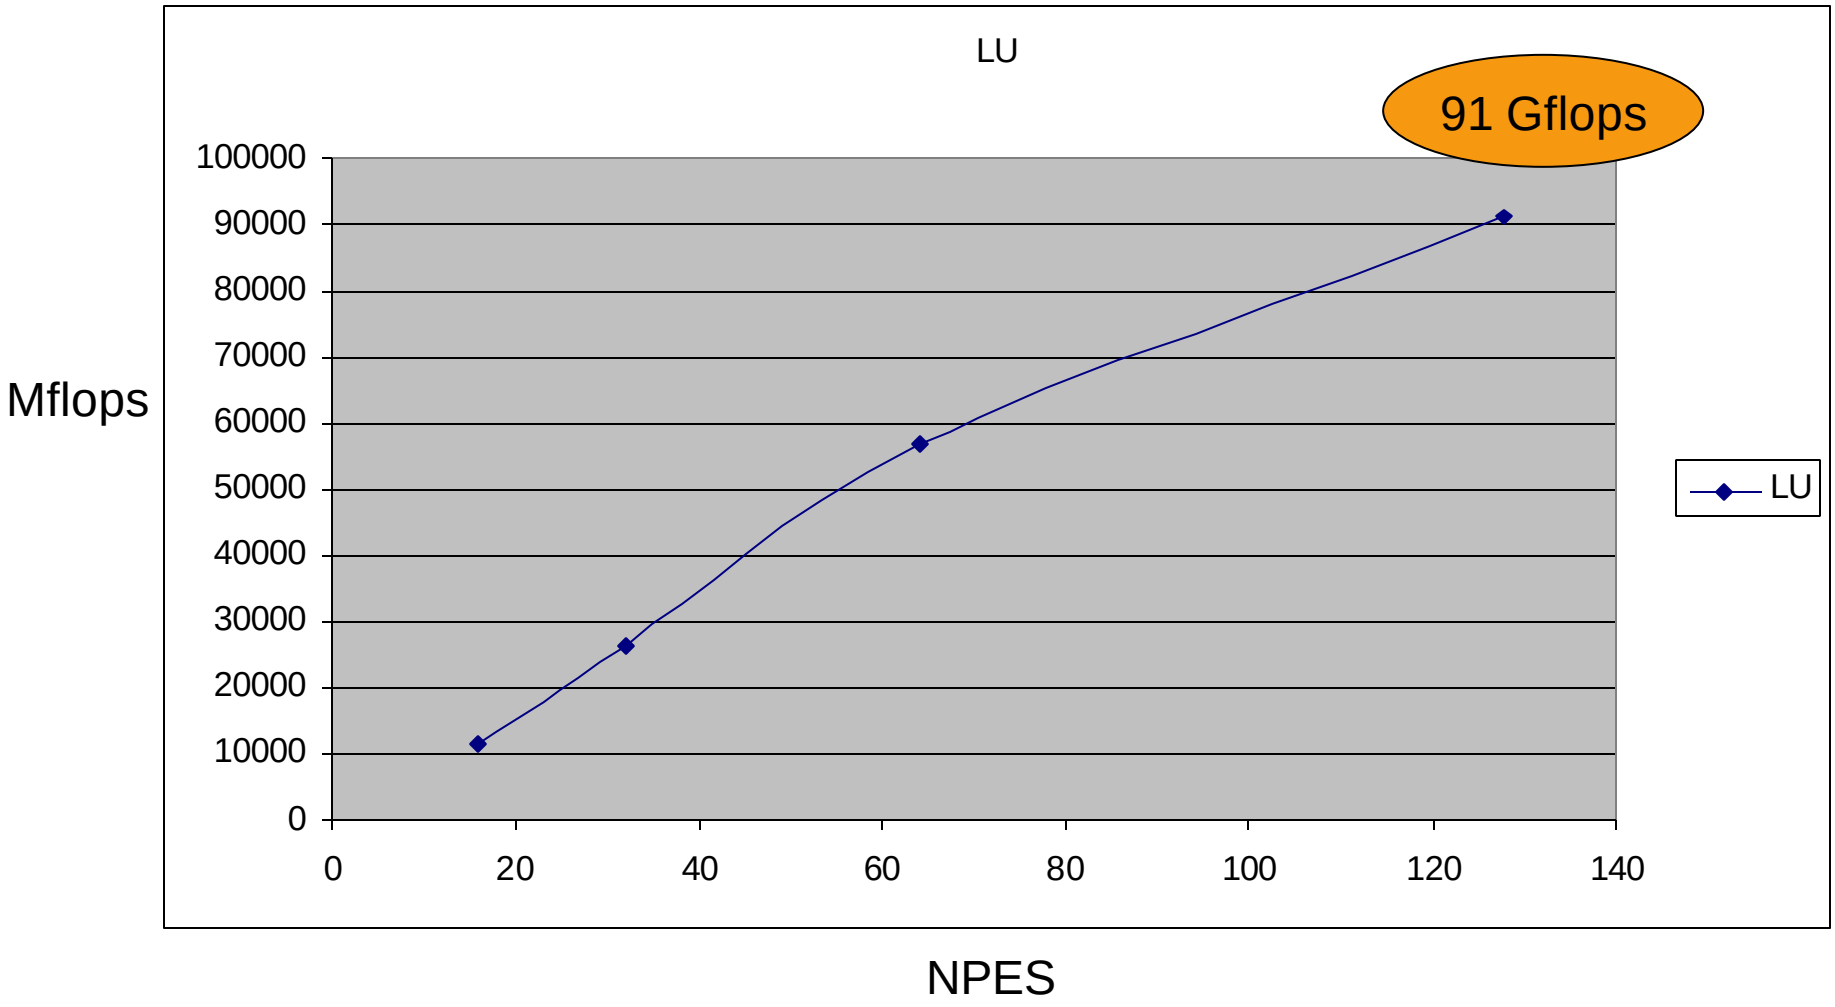
LU
91 Gflops
Mflops
100000
90000
80000
70000
60000
50000
40000
30000
20000
10000
0
0
20
40
60
80
100
120
140
LU
NPES

BT: Class C

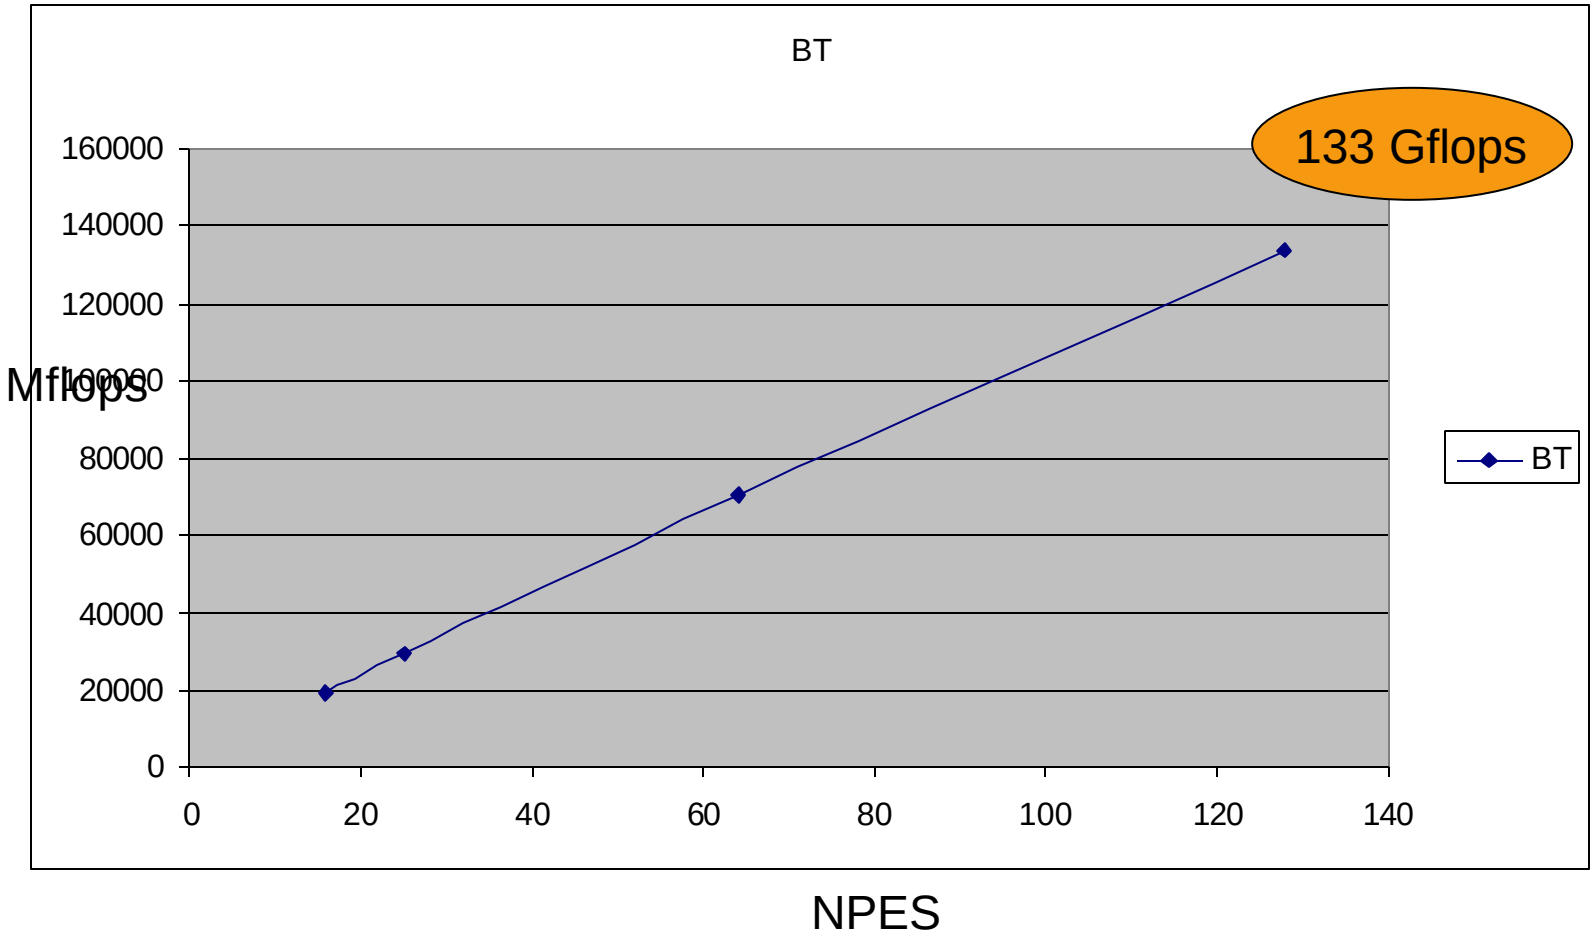
BT
133 Gflops
160000
140000
120000
100000
80000
60000
40000
20000
0
Mflops
0
20
40
60
80
100
120
140
NPES
BT

SP: Class C

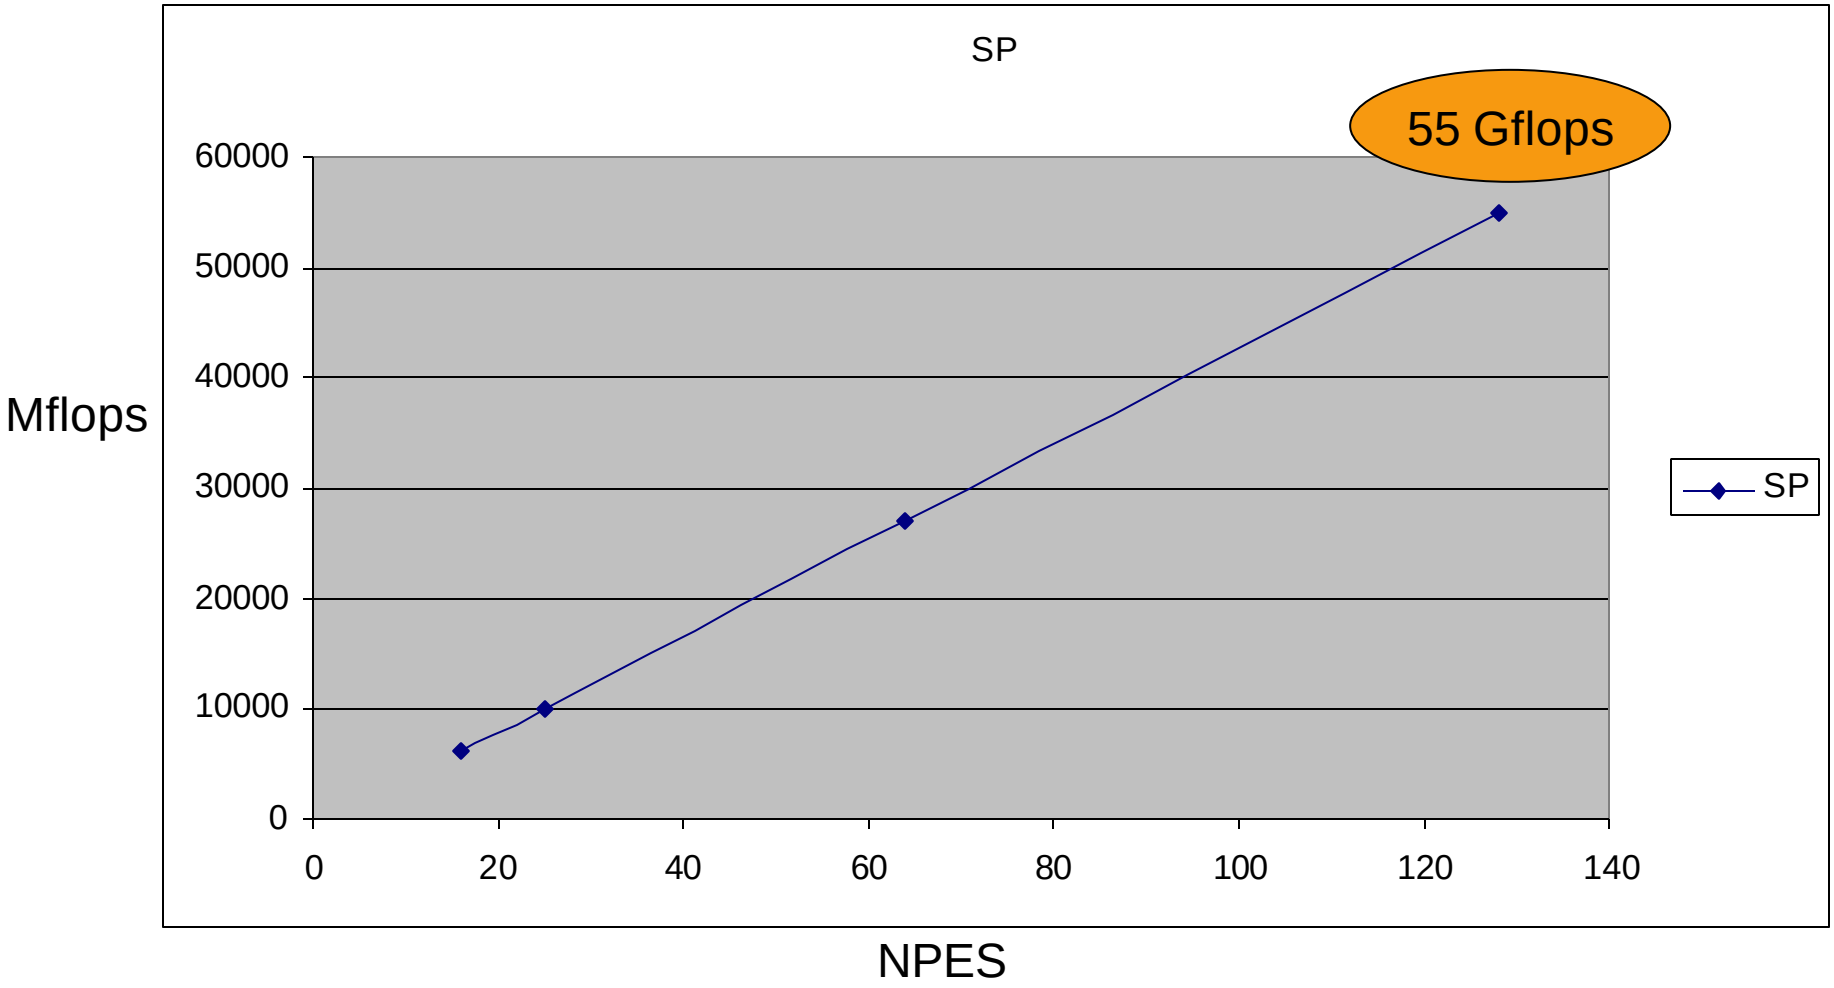
SP
55 Gflops
60000
50000
40000
30000
20000
10000
0
Mflops
0
20
40
60
80
100
120
140
NPES
SP

HPCC Ptrans Benchmark

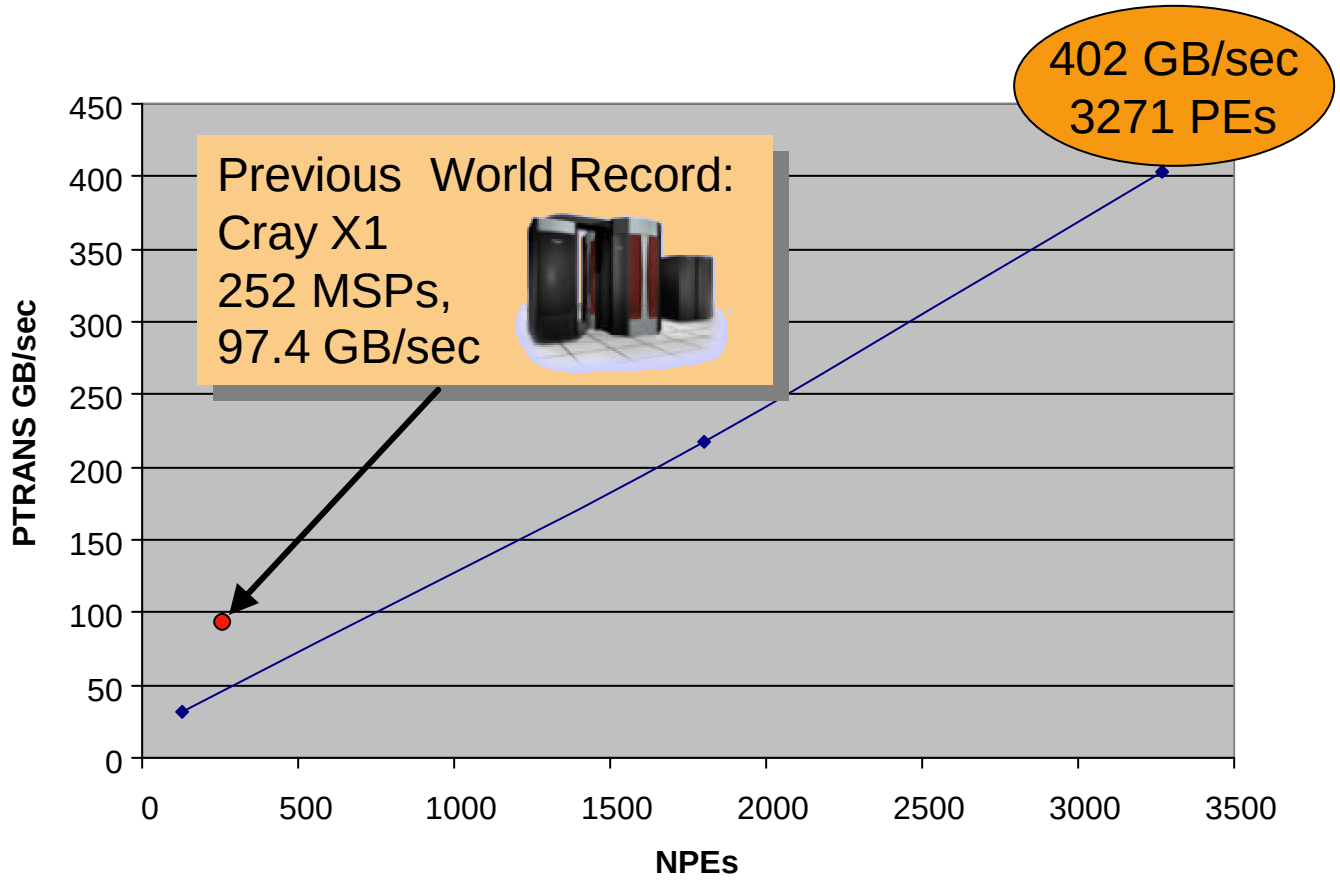

LM Weather Model

- 665 x 657 x 40 Grid
- 12 hour simulation, 40 second timestep

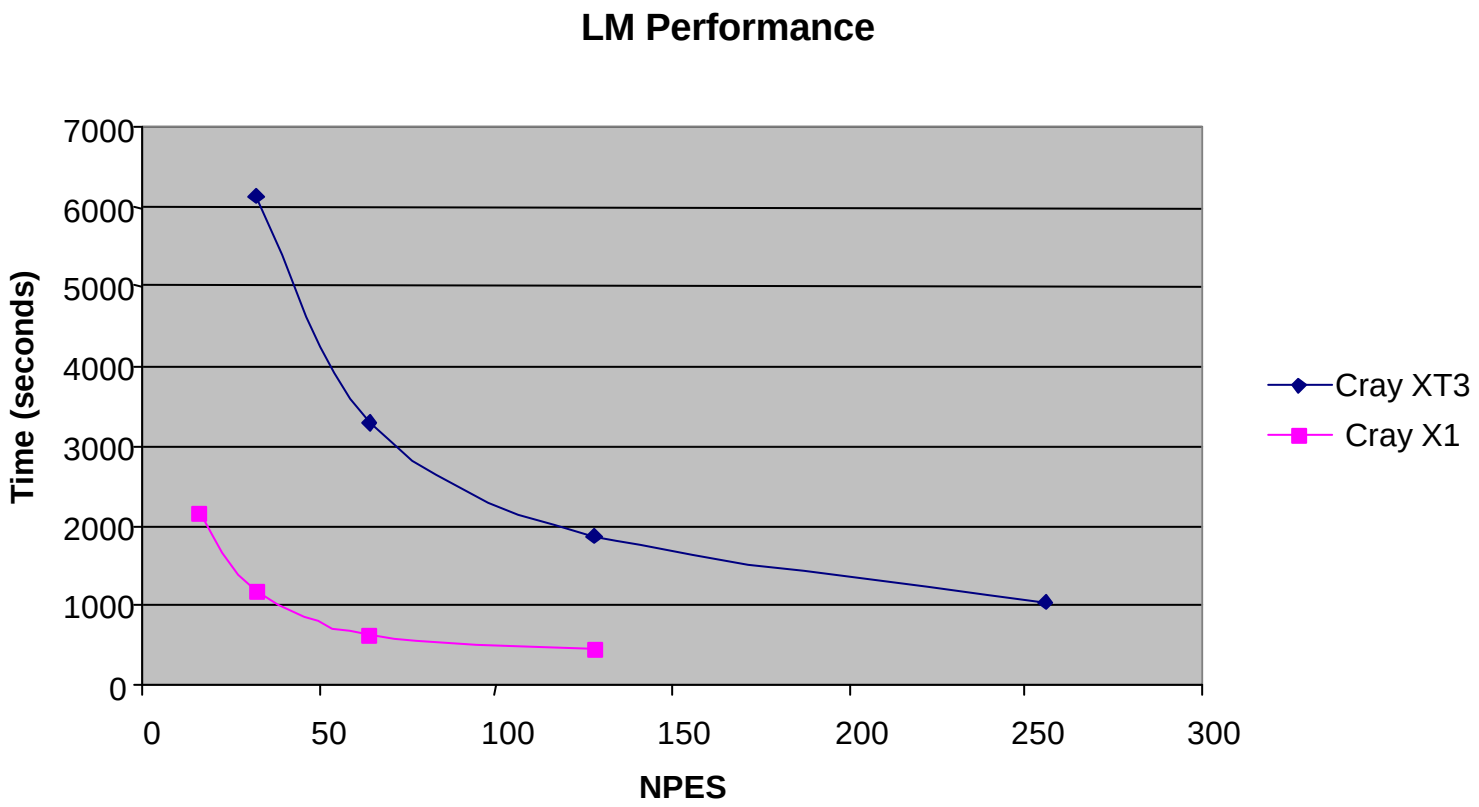

Codes Ported and running by Feb 2005:

- **Sandia 7x Apps**
  - Alegra
  - CTH
  - ITS
  - SAGE
  - Partisn
  - UMT2000
  - sPPM
  - Salinas
  - Presto
  - Calore
- **TI-05 Apps**
  - Aero
  - AVUS (Cobalt-60)
  - GAMESS
  - Hycom
  - RF CTH
  - WRF
  - Overflow
- **Research and Academic Chemistry**
  - Gromacs
  - NAMD
  - Amber 8
  - CPMD
- **Material Science**
  - LSMS
- **Weather/Climate**
  - ARPS
  - CAM
- **Other**
  - Quake
  - Gasoline
- **Benchmarks**
  - LINPACK
  - HPCC
  - NPB
  - STREAM
  - OSU Bi-section Bandwidth